The Effect of Distortions on the Prediction of Visual Attention

Milind S. Gide, Samuel F. Dodge, and Lina J. Karam, *Fellow, IEEE*

arXiv:1604.03882v1 [cs.CV] 13 Apr 2016

Abstract—Existing saliency models have been designed and evaluated for predicting the saliency in distortion-free images. However, in practice, the image quality is affected by a host of factors at several stages of the image processing pipeline such as acquisition, compression and transmission. Several studies have explored the effect of distortion on human visual attention; however, none of them have considered the performance of visual saliency models in the presence of distortion. Furthermore, given that one potential application of visual saliency prediction is to aid pooling of objective visual quality metrics, it is important to compare the performance of existing saliency models on distorted images. In this paper, we evaluate several state-of-the-art visual attention models over different databases consisting of distorted images with various types of distortions such as blur, noise and compression with varying levels of distortion severity. This paper also introduces new improved performance evaluation metrics that are shown to overcome shortcomings in existing performance metrics. We find that the performance of most models improves with moderate and high levels of distortions as compared to the near distortion-free case. In addition, model performance is also found to decrease with an increase in image complexity.

1 Introduction

Visual attention (VA) is the broad area of research that focuses on understanding the mechanisms by which the human visual system (HVS) filters out the huge amount of visual information collected by the retina. It enables humans to focus on the most relevant and important information that helps them perform complex cognitive tasks such as object recognition and scene interpretation in real-time. This process is thought to be a combination of two inter-dependent mechanisms, a bottom-up pre-attentive one, that relies on low-level data driven features and a top-down one that is cognitive, high-level and task dependent.

Over the past few decades, a lot of effort has gone into understanding the underlying bottom-up mechanisms of VA physiologically, psychologically as well as computationally. As a result, a large number of visual attention models (see [1] for a detailed review) have been proposed that mimic pre-attentive vision. Each of these approaches computes the saliency of a pixel giving a measure of how much that pixel stands out from its surroundings. Traditionally, these VA models are compared with eye-tracking fixations to gage their performance. A number of comprehensive comparative evaluation studies have been conducted in recent years [2], [3], to benchmark existing state-of-the-art models in terms of their ability to predict human fixations obtained from eye-tracking data. The ability of a VA model to select relevant areas in images is very useful for data selection and reduction in applications such as object recognition [4], [5] segmentation [6], [7] amongst others.

Another application of VA that has received considerable attention is image quality assessment (IQA). Images are subjected to a number of distortions at different stages of the image processing pipeline from acquisition to transmission. Thus, it becomes necessary to algorithmically detect the level of distortion affecting an image as perceived by a human observer. Several objective IQA algorithms have been proposed in the literature [8]–[10]. Most of these existing popular IQA algorithms pool information over all pixels to compute an image quality score. However, according to VA studies, humans tend to fixate only on certain salient areas in the image. Thus, pooling information only over perceptually important regions can improve performance. Based on this, a number of recent studies have explored the use of VA to improve IQA metrics [11]–[18]. Some of them [12]–[14] have used human eye-fixations to obtain saliency information that is used to weight the IQA metric scores at each pixel. Some others, [15]–[18], use VA models to improve the pooling of the IQA metrics. Except for one study [12], all others report a significant increase in IQA metric performance when combined with VA based pooling [11], [19], [20]. Some other studies [12], [21], [22] show that human fixations significantly change when the task given to subjects changes from free-viewing to quality assessment which is consistent with results reported in the well known study by Yarbus [23] which indicates significant changes in human scanpaths with a change in tasks. More recently, Mathe and Sminchisescu [24] also demonstrated the influence of different tasks on gaze under an object and context recognition framework. Given the potential of VA models in IQA algorithms in which the VA models

M. S. Gide, S. F. Dodge and L. J. Karam are with the School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, AZ, 85287-5706 USA (e-mail: mgide, sfdodge, karam@asu.edu)

are applied on distorted images, there are currently no evaluation studies in the literature that analyze the effect distortions have on the performance of state-of-the-art VA models. Also, given that almost all existing VA models have been developed for images without distortions, it is necessary to know how their performance changes due to distortions. This paper presents a comprehensive evaluation of 20 state-of-the-art VA models over distorted images obtained from two existing databases that provide eye-tracking data for a number of images with varying content and distortion types and levels. The change in performance over different levels and types of distortions is analyzed. As part of this evaluation, this paper also proposes new evaluation metrics that overcome shortcomings in existing performance metrics.

In this paper we provide the first comprehensive comparative evaluation of VA models on distorted images. We first give an overview of 20 state-of-the-art models that we chose to evaluate in Section 2. Next, in Section 3, we describe the performance metrics used to evaluate VA models and show shortcomings in existing performance metrics and propose new metrics that improve these shortcomings. Finally, in Section 4, we show and discuss the performance evaluation results and conclude the paper in Section 5.

2 Visual Attention Models

The VA models developed over the past few decades have mostly been targeted at modeling the bottom-up mechanisms of VA. This is because, top-down VA is highly dependent on the internal state of the observer and depends on factors like observer history which is very difficult to quantify and model. The bottom-up VA models work on the concept of a "saliency map" as proposed by Koch and Ullman [45]. The bottom-up VA models can be categorized based on the type of features used to obtain the saliency maps. We choose 20 state-of-the-art models that well represent different categories. The models also vary in complexity and correspondingly their resource intensiveness. The complete list of models along with the category, the time taken for each model to process a 768 x 512 image, the platform used, and the corresponding output resolution for each model is provided in Table 1. To calculate the computation time, all the models were run on the same PC using Matlab. For some models, MEX files and OpenCV libraries are used to speed up execution times, as indicated in Table 1 under the platform column. An overview of the different model categories and corresponding models is provided in Sections 2.1 to 2.6

2.1 Cognitive Models

Most of the VA models are based on cognitive principles directly or indirectly. However, some models are more heavily based on neurophysiological findings and give an insight into the neurological mechanisms of VA. One of the first basic models was developed by Itti et al. [26] based on the feature integration theory (FIT) [46]. In this model, the image is first decomposed into color, intensity and orientation channels. Then each of these channels is represented by Gaussian pyramids to obtain the center-surround responses that can be used to enhance features that are different from their neighbors. The different channels are combined across scales and features with normalization to give the saliency maps, the peaks of which correspond to visually more salient locations. We use two implementations of this model: ITTI1 refers to the original model proposed by Itti et al. [26] which is obtained via the implementation at http://www.saliencytoolbox.net, while ITTI2 refers to a higher resolution version of the same with Gaussian blurring and a center bias added to it, the implementation of which is found at http://www.klab.caltech.edu/~harel/share/gbvs.php. Another HVS based approach that models the mechanism of foveation is the Gaze Attentive Fixation Finding Engine (GAFFE) proposed by Rajashekar et al. [25]. GAFFE is a foveated analysis framework that uses four low-level local image saliency features, namely luminance, contrast, and bandpass outputs of both luminance and contrast, to form saliency maps and predict gaze fixations. It uses a sequential process in which the stimulus is foveated at the current fixation point and saliency features are obtained from patches from the foveated image to predict the next fixation point. The fixation points thus obtained can be converted to a saliency map by convolving them with a 2D Gaussian kernel with full-width-at-half-maximum equal to the visual angle corresponding to the foveal region.

2.2 Information Theory Based Models

These models are based on the principle that regions of high saliency maximize information and are described by rarity of features. Bruce et al. (AIM) [28] used Shannon's self-information measure [47] to get an indication of how unexpected the local image content is relative to its surroundings and saliency is taken to be proportional to the information content in a region in comparison with its surrounding. Zhang et al. (SUN) [32] used a similar approach by defining bottom-up saliency as the self-information of visual features; however, the self-information is described in terms of natural image statistics. Another similar saliency measure comprising of bottom-up and top-down components of saliency was proposed by Torralba (Torralba) [48] in which the bottom-up saliency is defined to be inversely proportional to the probability of features. Hou and Zhang (HouNIPS) [29] used the incremental coding length (ICL) concept to compute saliency wherein the objective is to maximize the entropy of the visual features. Mancas et al.

TABLE 1: List of Evaluated Models: W and H represent the width and height of the input image and are set to 768 and 512, respectively, for determining computation time, except when noted otherwise.

Category	Model Name	Time Taken (seconds)	Platform	Resolution
Cognitive	GAFFE [25]	5.5794	Matlab	$W \times H$
	ITTI1 [26]	0.4160	Matlab	$W \times H$
	ITTI2 [26]	4.9322	Matlab	$W \times H$
Graph Based	GBVS [27]	0.3515	Matlab + MEX	$W \times H$
	GBVSNoCB [27]	0.5496	Matlab + MEX	$W \times H$
Information Theoretic	AIM [28]	22.9508	Matlab	$\frac{W}{2} \times \frac{H}{2}$
	HouNIPS [29]	0.1960	Matlab	$W \times H$
	GR [30]	0.2926	Matlab	$W \times H$
	SDSRG [31]	1.2863	Matlab	$W \times H$
	SDSRL [31]	1.3263	Matlab	$W \times H$
	SUN [32]	2.8115	Matlab	246×331
	Torralba [33]	0.7909	Matlab	$W \times H$
Machine Learning Based	Judd [34]	298	Matlab + MEX+OpenCV	200×200
Spectral	FES [35]	0.3593	Matlab	$W \times H$
	FTS [36]	1.7581	Matlab	$W \times H$
	SigSal [37]	0.3852	Matlab	64×48
	SpectRes [38]	0.0862	Matlab	$W \times H$
Other	AWS [39]	5.1655	Matlab	$\frac{W}{2} \times \frac{H}{2}$
	BMS [40]	2.4805	Matlab + MEX	$W \times H$
	Context [41]	37.2912	Matlab + MEX	$W \times H$
	CovSal [42]	14.9193	Matlab	$W \times H$
	RandomCS [43]	3.4434	Matlab	$W \times H$
	SIMCoarse [44]	27.6926	Matlab + MEX + OpenCV	$W \times H$
	SIMFine [44]	28.7456	Matlab + MEX + OpenCV	$W \times H$

(GR) [30] also use self-information measures of the mean and variance of local image intensity, and the assumption that rare features attract visual attention. In the Saliency Detection by Self Resemblance (SDSR) approach proposed by Seo and Milanfar [31], local descriptors known as local regression kernels are computed over an image to measure the likeness or resemblance of a pixel with its surroundings. Saliency is then computed using this self-resemblance measure at each point, which measures the statistical likelihood of a saliency value given its neighborhood. The global (SDSRG) and local (SDSRL) variants of the model are evaluated separately. In the model proposed by Tavakoli et al. (FES) [35], a center-surround approach using a Bayesian framework is adopted. Saliency at a point is considered to be a binary random variable which is 1 if the point is salient and 0 otherwise. The probability of a pixel being salient given feature values is considered to be the saliency.

2.3 Spectral Analysis Based Models

Models described in this category compute saliency in the frequency domain. Hou and Zhang (SpectRes) [38] developed a method that models saliency by using the logarithm of the spectrum of an image and computing its deviation from the average log spectrum of several natural images, which acts as prior information. Achanta et al. (FTS) [36] developed a frequency tuned saliency approach in which an image is first converted to the CIE Lab color space and then low-pass filtered with binomial kernel to eliminate fine texture details and noise from which the mean of the image is subtracted to get the final high resolution saliency map with uniform salient regions. Another approach by Hou et al. [37] develops an image signature based on the sign function of the discrete cosine transform (DCT) which is used to predict visually conspicuous locations in an image.

2.4 Graph Based Models

A graphical model is a framework which describes the conditional dependence of neighboring elements or nodes. Harel et al. [27] introduced the Graph Based Visual Saliency (GBVS) model in which feature maps are extracted at multiple resolutions for different image features like intensity, color, and orientation. A fully connected graph, representing a Markov chain is constructed considering pixels of the feature maps as nodes. The weight between two nodes is assigned proportional to the similarity of feature values and spatial distance between them. On solving the Markov chain, the equilibrium distribution accumulates mass at nodes that have high dissimilarity with their surrounding nodes and the activation measure captures pairwise contrast. The GBVS model inherently promotes high saliency values near the center of the images. For a fairer comparison, we also include the GBVS method with the center bias removed for performance evaluation purposes. This is treated as a separate model denoted by GBVSNoCB in the results.

2.5 Machine Learning Based Models

Machine learning can be used to model VA by using eye-tracking data to train classifiers to predict salient vs. non-salient regions in an image. One such approach by Judd et. al (Judd) [34] uses several low-level features as input to a linear SVM classifier including

those in the original model by Itti et al. [26], Gabor filtered values and a global gist based saliency map along with higher level top-down features like face, horizon and object detectors. The machine learning problem is formulated as a linear regression in which the optimal weights for each feature are learned by training a linear regression model on a training set for which the eye-tracking data is used.

2.6 Other Models

Some of the newer approaches to VA are hybrid in the sense that they belong to more than one category or they cannot be readily classified into any of the aforementioned categories. The Adaptive Spectral Whitening (AWS) VA model proposed by Garcia-Diaz et al. [39] uses variability in local energy as a measure of saliency. The image is first transformed into the CIE Lab color space and the luminance channel is decomposed by both a multi-orientation as well as multi-resolution filter bank whereas the a and b channels are decomposed into a multi-resolution filter bank. These responses are first decorrelated using PCA and then locally averaged to obtain the saliency map. In another approach by Goferman et. al [41], called Context Based Saliency (Context), local as well as global distinctiveness is considered as a contributing factor towards saliency. The Saliency based on Image Manipulation (SIM) approach proposed by Margolin et al. [44] introduces an object probability map that groups together locations of high attention that are in close proximity to each other and also a pixel reciprocity measure that boosts the saliency of the pixels in close vicinity to highly salient locations. The coarse (SIMCoarse) and fine (SIMFine) versions of the model are evaluated separately. In the Boolean Map Saliency (BMS) approach proposed by Zhang and Scarloff [40], the Gestalt principle of surroundedness is used to compute saliency. In this approach, color-based feature maps are thresholded by variable thresholds to obtain boolean maps which are then used to obtain connected regions. Regions that exhibit surroundedness or are closed (i.e., regions not attached to the image boundary) are assigned a binary value of 1 and others are given a value of 0. Maps for a given threshold value are normalized using the L_2 norm and then averaged over different thresholds and different features to get the final saliency map. Another approach by Vikram et al. called Random Center Surround Saliency (RandomCS) [43] deviates from the Feature Integration theory based approaches [25], [26] by utilizing a stimulus-influence-based approach in which saliency at a point is influenced by other randomly chosen points in the image. The saliency influence at a point by another point is described in terms of a contrast by taking the absolute difference in pixel values and normalizing with the distance between the two points. This process is repeated over the L, a and b components of the image. The Covariance Saliency (CovSal) method proposed by Erdem et al. [42] uses covariance matrices of simple image features in the form of region covariance descriptors to capture local image structures and provide non-linear integration of the features. Saliency is obtained by finding the distance between the covariance matrices of a central region with its surrounding neighborhood regions using a non-Euclidean distance measure that is based on the eigenvalues and eigenvectors of the covariance matrices.

3 Performance Evaluation Metrics

The most common evaluation strategy is to analyze the performance of a VA model in terms of several metrics. However, existing saliency metrics will often give different rankings to different models (see [2]). It is unclear as to which metrics are more useful. In this section we describe existing metrics, propose four new metrics and show their advantage over the existing metrics.

We use the following common notation for all metrics. G is the ground truth saliency map (also called fixation density map), S is the saliency map produced by a model. (x_i,y_i) is the position of the ith ground-truth fixation point, where $0 \leq i \leq N-1$.

3.1 Existing Performance Metrics for VA models

3.1.1 Linear Cross Correlation (CC)

The linear cross correlation is a metric that measures the strength of a linear relationship between two variables. When used for evaluating VA models, it is used in the following form:

$$\mathrm{CC}(G,S) = \frac{\sum_{x,y}(G(x,y)-\mu_G)(S(x,y)-\mu_S)}{\sqrt{\sigma_G^2\sigma_S^2}} \quad (1)$$

where G represents the fixation density map obtained by convolving the fixation map (which is a map of 1s at fixation points and 0s at non-fixated points) with a 2D Gaussian, and S represents the predicted saliency map. μ_S, μ_G are the means and σ_G, σ_S are the standard deviations of the predicted saliency map and fixation density map, respectively. The CC value can vary between -1 and 1, with a value of 1 denoting a strong positive linear relationship, a value of -1 representing a strong negative linear relationship and a value of 0 denoting absence of correlation.

3.1.2 Similarity (SIM)

The similarity measure was presented by Judd et al. [49] as part of a benchmark for VA models. It corresponds to the histogram intersection between the predicted saliency map histogram H_S and the fixation density map histogram H_G. It is defined as follows:

$$\mathrm{SIM}(S,G) = \sum_{i=1}^{N} min(H_S(i), H_G(i)); \quad (2)$$

where N is the number of histogram bins.

3.1.3 Earth's Mover Distance (EMD)

The Earth Mover's Distance (EMD) [50] is a cross-bin histogram dissimilarity measure that doesn't require the histogram domains to be aligned. The EMD treats two histograms or distributions as two different ways in which the same amount of dirt has piled up over a region, and the distance is given by the minimum cost of changing one pile of dirt into the other. An improved variant of the EMD metric proposed by Pele and Werman [51] (denoted by $\widehat{\text{EMD}}$), that works as a metric even for histograms that are not normalized, is given by:

$$\widehat{\text{EMD}}(H_S, H_G) = \left(\min_{\{f_{i,j}\}} \sum_{i,j} f_{i,j} d_{i,j} \right) + |\sum_i H_G(i) - \sum_j H_S(j)| \max_{i,j} d_{i,j} \quad (3)$$

$$\text{s.t } f_{i,j} \geq 0,\ \sum_j f_{i,j} \leq H_G(i),\ \sum_i f_{i,j} \leq H_S(j),$$

$$\sum_{i,j} f_{i,j} = \min(\sum_i H_G(i), \sum_j H_S(j))$$

where H_S and H_G represent the histograms of the predicted and ground-truth saliency maps, respectively, $f_{i,j}$ represents the amount transported from bin i to bin j, and $d_{i,j}$ is the ground distance between bin i and bin j in the histograms. For probability histograms (total mass of 1), the EMD and $\widehat{\text{EMD}}$ values are equal. $\widehat{\text{EMD}}$ was used by Judd et al. in their VA benchmark [49] with the ground distance saturated to a maximum value as proposed in [52].

3.1.4 Area under the Receiver Operating Characteristic Curve using Ground Truth Fixations (AUC-F)

For this evaluation method which is based on the receiver operating characteristics (ROC), the predicted saliency map is first thresholded. This thresholded map is denoted by S_t, where t represents the threshold value which ranges from 0 to 255 for an 8-bit depth saliency map.

Then, the true positive rate (TPR) and false positive rate (FPR) are determined using the thresholded map and the ground-truth fixation points. The true positive rate is computed by counting the fraction of ground-truth fixation points that are properly detected in the thresholded saliency map as follows:

$$\text{TPR(t)} = \frac{1}{N} \sum_{i=0}^{N-1} S_t(x_i, y_i) \quad (4)$$

To find the FPR we must compute the fraction of non-fixated points that are labeled as salient in the thresholded predicted saliency map. One approach [2] for computing the FPR is to have the number of non-fixated points be equal to the number of fixated points. N random points are chosen from the possible set of non-fixated pixels in the ground-truth. We use the notation $(\hat{x}_i, \hat{y}_i)$ to denote these non fixated points. Because these points are chosen randomly, the AUC test is performed 100 times each with different random non-fixated points, and the final result is the average of the individual trials. Let l be the index of the current trial and L be the total number of trials. Also, let $(\hat{x}_{i,l}, \hat{y}_{i,l})$, $1 \leq i \leq N$ be the set of N non-fixated points at the l^{th} trial. The FPR is then given by:

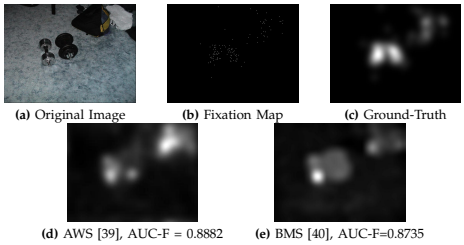

(a) Original Image **(b)** Fixation Map **(c)** Ground-Truth

(d) AWS [39], AUC-F = 0.8882 **(e)** BMS [40], AUC-F=0.8735

Fig. 1: Disadvantage of AUC-F in terms of evaluating a model's ability to detect the most salient regions.

$$\text{FPR}(t) = \frac{1}{L} \sum_{l=1}^{L} \text{FPR}(t, l)\,, \quad (5)$$

where

$$\text{FPR}(t, l) = \frac{1}{N} \sum_{i=1}^{N} S_t(\hat{x}_{i,l}, \hat{y}_{i,l}) \quad (6)$$

By computing the *TPR* and *FPR* for every value of the threshold t (from 0 to 255) we obtain a receiver operating characteristic curve (ROC). The area under this curve (AUC) gives a measure of the performance of the saliency map. A value of 1 means that the saliency map perfectly predicts the ground-truth, whereas a value of 0.5 is no better than random chance.

3.1.5 Area under the Receiver Operating Characteristic Curve using Ground Truth Saliency (AUC-S)

The AUC-F metric gives equal importance to each ground truth fixation. As a result, the highly salient regions that attract more fixations per unit area in an image are treated the same as other less salient regions that attract fewer fixations per unit area. Figure 1 shows an example from the Toronto database [28] that illustrates this issue. As seen from the fixation map and the ground truth saliency map, which is obtained by placing 2D Gaussians with full-width-at-half-maximum (FWHM) corresponding to one degree of visual angle on each of the fixation points, the area with the dumbbells has a higher density of fixations as compared to the area with the bag. However, AUC-F will treat fixations near the dumbbells the same as those near the bag. As a consequence, some VA models whose saliency map incorrectly predicts the bag to be a more salient object than the dumbbells (e.g., AWS [39] in Figure 1(d)) will get a higher AUC-F score

than VA models (e.g., BMS [40] in Figure 1(e)) whose saliency map is more in agreement with ground truth data. A saliency map variant of the AUC (denoted by AUC-S) that uses the ground truth saliency maps instead of fixations attempts to rectify this problem. For this metric, first, a fixed threshold $T = 0.5 * \sigma$, where σ is the standard deviation of G, is applied to G to get a binary map $G_T = \{0, 1\}$. Then, a variable threshold t is applied to S to get a binary map $S_t \in \{0, 1\}$. Let $I_t = G_T \wedge S_t$ denote the intersection of both binary maps, $n(X)$ the number of non-zero points in a binary map X, and $l(X)$ the total number of points in X. The TPR and FPR can then be computed as follows:

$$\text{TPR(t)} = n(I_t)/n(G_T) \tag{7}$$

$$\text{FPR(t)} = \frac{n(S_t) - n(I_t)}{l(G_T) - n(G_t)} \tag{8}$$

The TPR and FPR are computed for every value of the threshold t to obtain a receiver operating characteristic curve (ROC) the area under which gives AUC-S. Like the AUC-F, for AUC-S, a value of 1 indicates the best prediction performance, whereas a value of 0.5 indicates chance performance.

3.1.6 Shuffled AUC (SAUC)

The main limitation of the AUC-F and AUC-S metrics is that they are center biased. When using these metrics, a model's score can be improved simply by multiplying the predicted saliency map by a large Gaussian blob such that the predicted saliency near the edge of the image will be reduced [32]. In fact, a simple Gaussian blob will often perform better than some saliency models using the AUC-F and AUC-S metrics. This center bias is inherent to eye tracking experiments. This occurs for two main reasons. First, the most salient regions of an image are often located near the center of the image. This is called the *photographer's bias*. Secondly, the nature of eye tracking experiments yields a framing effect, where it is natural to look at the center of the scene because this point allows our vision system to capture the most information [53].

The shuffled AUC (SAUC) metric was introduced to account for this center bias [32]. An unbiased metric should give a low score to the central Gaussian blob model. The difference in the SAUC metric as compared to the AUC metric is that the non-fixated points $(\hat{x}_i, \hat{y}_i)$ are randomly chosen from the distribution of all the fixation points from all the other images in the dataset. Consequently, this will give the Gaussian blob model an SAUC of approximately 0.5, which is equivalent to chance. The center-bias issue for AUC-F and AUC-S is illustrated in Figure 2 which shows the AUC-S, AUC-F, CC and SAUC scores for a centered Gaussian Blob in comparison with the AIM model [28] for an image taken from the Toronto database [28]. The AUC-F, AUC-S and CC scores for the centered Gaussian are higher than that for the AIM model because of the center bias. The SAUC score results, for the centered Gaussian, in a score that is closer to 0.5 and that is lower than the AIM model's score.

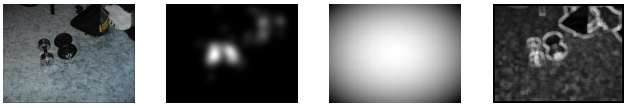

(a) Original Image **(b)** Ground Truth **(c)** Centered Gaussian **(d)** AIM [28]

	(c)	(d)
AUC-F [2]	0.8079	0.7551
AUC-S [54]	0.7991	0.7207
CC [49]	0.4399	0.3747
SAUC [2]	0.5823	0.6679

Fig. 2: Center bias problem in existing metrics that is rectified by the shuffled metrics like SAUC.

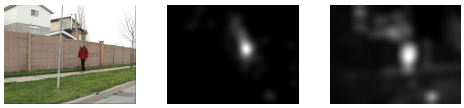

(a) Original Image **(b)** Ground-Truth **(c)** BMS [40]

	(b)	(c)
SAUC [2]	0.8052	0.8234
NSS [55]	4.1829	2.7881

Fig. 3: Image from Toronto [28] database that shows the advantage of the NSS metric over the SAUC metric. In terms of SAUC, the BMS model is incorrectly given a higher score than the ground truth saliency map.

3.1.7 Normalized Scanpath Saliency (NSS)

The normalized scanpath saliency (NSS) computes a score using ground-truth fixation points and the predicted saliency map as follows:

$$\text{NSS(x,y)} = \frac{1}{N} \sum_{i=1}^{N} \frac{S(x_i, y_i) - \mu_s}{\sigma_s} \tag{9}$$

where S is the saliency map to be tested, (x_i, y_i) is the location of a fixation in the ground-truth eye tracking data, μ_s and σ_s are, respectively, the mean and standard deviation of the predicted saliency map.

Compared to the NSS, the SAUC score, previously defined, suffers from the interpolation flaw (see [56] for more details) as a result of which models that output maps that are blurry and fuzzy tend to do better in terms of the SAUC metric. This can be observed in images with single, small salient regions and ground truth saliency maps that are sparse. One such image and its corresponding ground truth saliency map is shown in Figure 3 along with the SAUC scores for the ground truth saliency map and a state-of-the-art model BMS [40]. The BMS model is incorrectly given a higher SAUC score than the ground truth saliency map as true positives are given more importance than false positives in the SAUC formulation. The NSS does not have this bias and has a penalty for false positives. Hence it is a more fair metric to analyze models in the case of sparse saliency maps. However, NSS suffers from the same problem as the previously introduced AUC metrics (AUC-F and AUC-S) in that

(a) Ground Truth **(b)** Predicted Map S using BMS [40], SKLD [32] = 9.36 Proposed SSKLD = 9.36 **(c)** Inverted Map SKLD [32] = 9.46 Proposed SSKLD = -9.46

Fig. 4: Limitation of shuffled KLD from [32]. The original SKLD metric gives good scores for inverted saliency maps (c). Our proposed modified metric (SSKLD) gives a negative score for this inverted map.

it is sensitive to center bias. To counter this center-bias we introduce a new modified NSS metric termed shuffled NSS (SNSS). The proposed SNSS is described in more details in Section 3.2.

3.1.8 Shuffled Kullback-Leibler divergence

The Kullback-Leibler Divergence (KLD) [57] is widely used as a distance measure between two histograms. A shuffled KLD (SKLD) was used in [32] to compare histograms of the fixated and non-fixated points. As in the case of SAUC, the non-fixated points were chosen randomly from the distribution of fixation points of all other images in the dataset. If the histograms are very different then the SKLD score should be high and the saliency map is considered a good predictor of the ground-truth.

The motivation behind including this shuffled KLD (SKLD) score in addition to the other metrics is described in [58]. The authors of [58] show that out of all other metrics the SKLD measure is the most decorrelated with the AUC and NSS measures and hence provides a non-redundant metric of performance.

However, the SKLD metric presented in [32] is flawed because the metric compares the distance between these histograms regardless of which histogram has a larger average value. Because of this, if the gray-level values of a predicted saliency are inverted, the SKLD metric will give the same score (Figure 4). Such inverted saliency maps should however be given a lower score as they predict the non-salient regions as salient instead of the salient regions.

We propose a modification of the SKLD metric that does not exhibit this problem. The proposed metric is referred to as signed shuffled KLD or SSKLD for short and is described in Section 3.2 . This metric can give negative scores if the mean of the distribution of the non-fixated points is higher than that of the fixated points.

The KLD based measures, SKLD and SSKLD, cannot be classified as distance metrics as they do not have an upper bound and do not satisfy the triangular inequality. To resolve this issue we propose a performance metric based on the square root of the symmetric and bounded variant of the KLD known as the Jensen-Shannon Divergence (JSD) [59] which can be considered as a distance measure. The proposed metric termed as the Shuffled Jensen-Shannon Distance (SJSD) is described in Section 3.2.

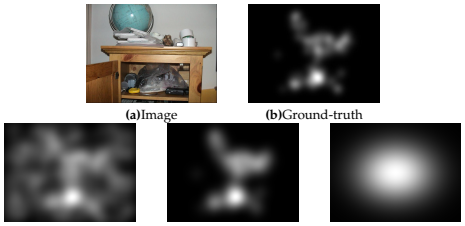

(a) Image **(b)**Ground-truth **(c)** **(d)** **(e)**

	(c)	(d)	(e)
AUC [60]	0.9056	0.9190	0.7993
SAUC [2]	0.8145	0.8120	0.4603
NSS [55]	1.9630	2.2598	1.0790
Proposed SNSS	1.3557	1.6211	-0.1400

Fig. 5: Advantage of the new SNSS metric. The two saliency maps (c) and (d) have nearly identical SAUC scores, however it is clear that (d) is a "better" saliency map. The proposed SNSS metric that we introduce in this paper is a better measure of saliency prediction. Here (d) has a much higher SNSS than (c). Additionally, a centered Gaussian blob (e) will perform well on the NSS score, however using the proposed SNSS the same Gaussian blob receives a score close to 0.

3.2 Proposed Performance Metrics for Evaluation Models

The existing performance metrics described previously have either one or both of the following two shortcomings, center bias and false positives. Center bias is related to the tendency of some metrics to unfairly reward higher central activity and false positives relates to the tendency of the AUC metrics (AUC-F, AUC-S and SAUC) to not penalize false-positives which gives an unfairly high score to blurrier maps. To resolve these issues we propose the following performance metrics.

3.2.1 Shuffled Normalized Scanpath Saliency

The proposed shuffled NSS (SNSS) metric is based on the NSS [55] metric and is defined as follows:

$$\text{SNSS} = \text{NSS}(x, y) - \text{NSS}(\hat{x}, \hat{y}) \tag{10}$$

where (x_i, y_i) and $(\hat{x}_i, \hat{y}_i)$ are, respectively, the ground-truth fixation and non-fixated points for the considered image. As before, the non-fixated points are chosen randomly from the fixation points of all other images in the dataset. This computation is repeated 100 times and the final result is the average of these 100 trials. Figure 5 shows the advantage of the SNSS metric over SAUC by creating two different synthetic saliency maps that are obtained by adding high and low amounts of background activity to the ground truth saliency map. The SAUC gives a higher score to the map with higher level of background activity whereas the SNSS correctly gives the map with lower background activity the higher score. SNSS also retains the invariance to center-bias that the SAUC has.

3.2.2 Signed Shuffled Kullback Leibler Divergence (SSKLD)

The SSKLD is obtained by modifying the shuffled KLD measure by multiplying it with the sign of the SNSS value defined in (10). Mathematically,

$$\text{SSKLD} = sign\Big(\text{SNSS}\,(S)\Big) \cdot \text{KLD}\Big(H, \hat{H}\Big) \tag{11}$$

where H is a histogram formed from (x_i, y_i) and $\hat{H}$ is a histogram formed from $(\hat{x}_i, \hat{y}_i)$. Both H and $\hat{H}$ are normalized by dividing the bin frequencies with the total number of ground truth fixations. As indicated previously, $(\hat{x}_i, \hat{y}_i)$ are chosen randomly from the ground-truth fixation points of all the other images in the dataset. KLD is the symmetric KLD distance given by:

$$\text{KLD}(H, \hat{H}) = \frac{\sum_i H(i) \log\left(\frac{H(i)}{\hat{H}(i)}\right) + \sum_i \hat{H}(i) \log\left(\frac{\hat{H}(i)}{H(i)}\right)}{2} \tag{12}$$

In practice, a small ϵ is added to $\hat{H}$ and H to avoid a division by zero and $\log(0)$ errors.

As seen in Figure 4, the signed term in the proposed SSKLD metric enables it to penalize the inverted saliency map by giving it a negative score, whereas using the SKLD metric results in the same score for both the inverted and the non-inverted maps.

3.2.3 Shuffled Jensen-Shannon Distance (SJSD)

The Jensen-Shannon Divergence between two histograms P and Q is a symmetric divergence based on the KLD and is given by

$$\text{JSD}(P, Q) = \frac{D(P||M) + D(Q||M)}{2} \tag{13}$$

where $M = \frac{P+Q}{2}$ is the mean of the two histograms and $D(P||Q) = \sum_i P(i) \log_2\left(\frac{P(i)}{Q(i)}\right)$ is the non-symmetric Kullback-Leibler Divergence. When a logarithm of base 2 is used for the calculation of $D(P||Q)$, $0 \leq \text{JSD} \leq 1$. The square root of the JSD is a distance metric that satisfies the triangle inequality [61]. The proposed shuffled Jensen-Shannon Distance (SJSD) is defined as follows:

$$\text{SJSD}(H, \hat{H}) = \Big(\text{JSD}(H, \hat{H})\Big)^{\frac{1}{2}} \tag{14}$$

where H and $\hat{H}$ are the same as in (11).

The SJSD measure improves on the SKLD and SAUC measures by penalizing false-positives as shown in Figure 6. It can be seen that SAUC and SKLD both incorrectly give the saliency map for the BMS [40] a higher score than the ground truth fixation density map which is rectified by the proposed SJSD metric.

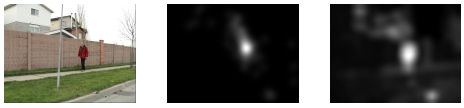

(a) Original Image (b) Ground Truth (c) BMS [40]

	(b)	(c)
SAUC [2]	0.8052	0.8234
SKLD [32]	0.2138	0.4064
Proposed SJSD	0.8278	0.6417

Fig. 6: Image from Toronto [28] database that shows the advantage of the proposed SJSD metric over the SAUC and SKLD metrics. In terms of SAUC and SKLD, the BMS model is incorrectly given a higher score than the ground-truth saliency map. The proposed SJSD metric gives a higher score to the ground-truth saliency map as expected.

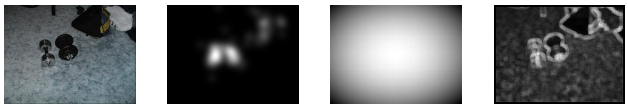

(a) Original Image (b) Ground Truth (c) Centered Gaussian (d) AIM [28]

	(c)	(d)
$\widehat{\text{EMD}}$ (lower is better) [49]	2.4060	2.7490
Proposed SEMD (higher is better)	1.0185	1.2789

Fig. 7: Center bias problem in the $\widehat{\text{EMD}}$ [49] metric that is rectified by the proposed SEMD metric.

3.2.4 Shuffled Earth Mover's Distance (SEMD)

As seen in Figure 7, the disadvantage of using the $\widehat{\text{EMD}}$ metric defined in (3) to compare the histograms of the ground-truth and predicted saliency maps is that it is center-biased as it incorrectly assigns the centered Gaussian blob a higher score than the VA model. To rectify this issue, we propose a shuffled version of the $\widehat{\text{EMD}}$ metric that is defined as follows:

$$\text{SEMD} = \widehat{\text{EMD}}(H, \hat{H}) \tag{15}$$

where H and $\hat{H}$ are the same as in (11). Contrary to the original $\widehat{\text{EMD}}$ of (3) (lower score means better), a higher SEMD score indicates better performance. The shuffled EMD score correctly assigns a lower score to the centered Gaussian blob as compared to the VA model. We use a fast implementation of the $\widehat{\text{EMD}}$ metric proposed in [52] (code provided at http://www.cs.huji.ac.il/~ofirpele/FastEMD/code/) that makes use of a saturated ground distance in computing $\widehat{\text{EMD}}$.

4 Performance Evaluation

In order to evaluate the effect of different types and levels of distortions on the performance of VA models, predicted saliency maps are obtained for images that have been subjected to different levels and varying types of distortions. These are then compared with corresponding eye-tracking data results for which we use the proposed shuffled SNSS, SSKLD, SJSD and SEMD metrics introduced in Section 3.2 along with the widely used SAUC metric. As the different models produce saliency maps of differing resolutions (see Table 1), we resize the saliency map given by each model to the size of the original image before

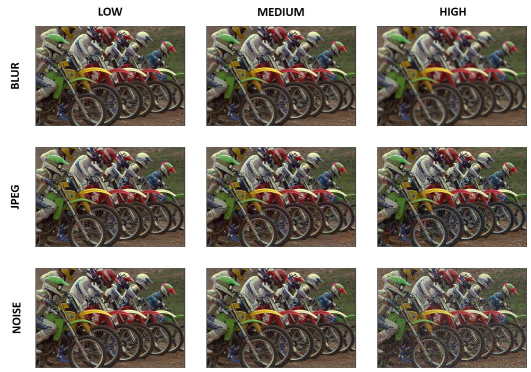

Fig. 8: Distorted versions of the Bikes image in the TUD database [62]. The level of distortion increases from left to right and the type of distortion changes from top to bottom, namely, Blur, JPEG and Gaussian Noise, respectively.

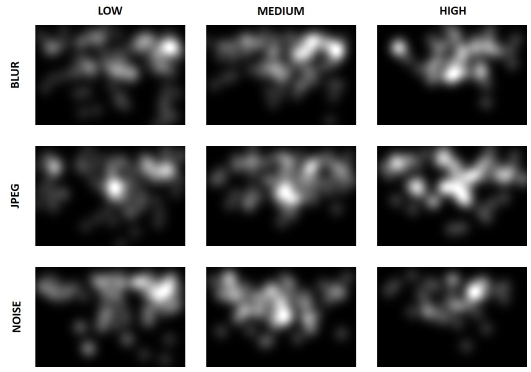

Fig. 9: Fixation density maps for distorted versions of the Bikes image in the TUD database [62]. The level of distortion increases from left to right and the type of distortion changes from top to bottom.

computing the performance metric score as in [2]. Also, as blurring the saliency map improves scores over certain metrics, we blur the predicted saliency map over different blur levels and choose the optimal blur level that maximizes the score as in [2], [40]. For comparing the models with each other, we choose this maximum value at the optimum blur level.

4.1 Databases

4.1.1 TUD Interactions Database

The TU Delft Interactions Database [21] is used as the ground-truth eye-tracking data for evaluating the performance over distorted images. This database includes human eye movements recorded for 14 subjects while looking at 54 distorted stimuli. The different distortion types used are Gaussian blur, white noise and JPEG compression, with each type having three different levels of distortion (high, medium and low). The database provides MOS scores from the quality assessment task in addition to the saliency maps obtained from the recorded fixation points. Figure 8 shows the distorted versions (different types and different levels of distortion) for the Bikes image, and Figure 9 shows the corresponding fixation density maps for the images. The fixation density maps are

Fig. 10: Sample images covering all resolutions from the Judd Low Resolution Database [63] for each of the three categories: Easy, Medium and Hard.

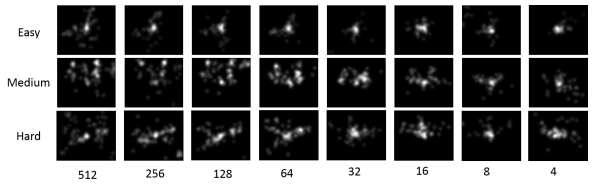

Fig. 11: Fixation density maps for images shown in Figure 10 covering all resolutions from the Judd Low Resolution Database [63] for each of the three categories: Easy, Medium and Hard.

obtained by placing a 2D Gaussian on the fixation point locations and normalizing the resulting map.

4.1.2 Judd Low Resolution Database

The Judd Low Resolution database [63] contains eye tracking data for original versions of 168 natural images and 25 pink noise images in addition to down-sampled versions of the images varying from 4 to 512 in terms of image height. The images are shown to 64 observers after scaling them to a size of 860×1024. The natural images are divided into three categories, easy, medium and hard, depending on image complexity. The categorization is done using a subjective criteria in which first, each image is assigned a score that corresponds to the lowest resolution at which the image is understood by a subject. The images are then ranked in the ascending order and binned into three equal sized categories. The images with the lowest scores make up the easy category and those with the highest scores form the hard category with the remaining images making up the medium category. Based on the eye tracking data, it is observed that for the easy images, fixations stay consistent even for extremely low resolutions and are more centered. On the other hand, the hard images have more variations in eye fixations which are less centered for higher resolutions because there are more details on which subjects fixate. As the input images are subsampled at different levels, the distortion that each image is affected with, in this case, is equivalent to blur. The lower the resolution of the image, the higher is the level of blur.

Figure 10 shows a sample image from each of the three categories for the different resolutions and Figure 11 shows the corresponding fixation density maps that are obtained by convolving a 2D Gaussian kernel over all the fixations for the images and then normalizing the resulting maps. In our analysis, we

TABLE 2: Consistency of performance for existing and proposed metrics over the MIT [34] and Toronto [28] datasets in terms of the Kendall's coefficient of concordance (KCC).

	SAUC	SNSS	SSKLD	SEMD	SJSD
KCC	0.8978	**0.9265**	0.7778	0.8704	0.8678

TABLE 3: Top 5 Models over 3 or more metrics for the TUD database, for different levels and types of distortions.

Blur			JPEG			Noise		
High	Medium	Low	High	Medium	Low	High	Medium	Low
ITTI	SDSRL	SDSRG	RandomCS	SimCoarse	SDSRG	SCIA	Judd	Judd
SCIA			GR	SDSRG	SDSRL	SDSRL	RandomCS	
Judd			SDSRG			CovSal	SCIA	
CovSal			Context				CovSal	

only consider the natural images (1344 out of 1544) in the dataset and ignore the colored noise images because the considered saliency models give a meaningful output for images that have some inherent structure and which are not random in nature.

4.2 Results

4.2.1 Performance over non-distorted images

We first validate the proposed metrics by using them to evaluate the VA models listed in Table 1 on the widely used Toronto [28] and MIT [34] databases which consist of non-distorted images. Figure 12 shows the performance for the proposed metrics as well as the SAUC metric [2]. To validate the consistency of the metrics over the two datasets, we adopt as in [58] the Kendall's coefficient of concordance (KCC) to compare the models' rankings provided by each metric over the two datasets. Table 2 shows the obtained results. From Table 2, it can be seen that the proposed metrics provide consistent ratings in terms of models' rankings over the two datasets with the proposed SNSS metric showing the best consistency.

4.2.2 Variation in performance with varying distortion types and levels for the TUD Interactions database

As indicated previously, the TUD Interactions database has three different types of distortions: Gaussian blur, JPEG compression and Gaussian noise. Each of these distortion types are present in three different levels, high, medium and low. To check the performance over different distortion types, as well as the effect of the different levels of distortions have on the predicted saliency given by different models, we compute metric scores for each image and then average them over those set of images that belong to a particular distortion type and a particular level of distortion.

Figure 13 shows the performance of each model for the different distortion levels for a particular distortion type in terms of the SAUC, SNSS, SSKLD, SJSD and SEMD metrics, respectively. In each case the models are sorted such that the average peformance over different levels of distortions increases from left to right. Table 3 shows the models that occur in the top five for three or more metrics. Tables 4 and 5 show the average standard deviation of the metric scores for all the models over different levels and types of distortion, respectively. To account for the different ranges of scores for each of the metrics, the metric scores obtained for every model are normalized by the maximum value of the metric over all distortion types for each distortion level (Table 4) or over all distortion levels for each distortion type (Table 5) before computing the standard deviation. We can draw the following conclusions from Figure 13 and Tables 3 to 5:

- The performance for each model changes with a change in distortion type as the ranking of the models in terms of average performance over different levels of distortions changes with the type of distortion for a given metric. Also, the performance for each model varies with a change in distortion levels. It is found that increasing the level of distortion improves performance to a great extent for the blur type of distortion.
- The performance scores for the different metrics show that performance is higher for the high and medium levels of distortions than low levels of distortions (near distortion-free) for most of the models over all metrics as well as all distortion types. This result suggests that addition of certain levels of distortion helps the VA models predict the salient areas in the distorted image. As the eye-tracking results for this database have been recorded for a quality assessment task, it is highly probable that subjects fixated upon distortions which the saliency models capture thereby helping with metric performance.
- If we consider the performance of models over different metrics, as shown in Table 3, SDSR [31] performs the best overall by ranking in the top five for three or more of the five metrics for the medium and low distortion levels for the blur distortion type, the high, medium and low distortion levels for the JPEG distortion type, and the high distortion level for the noise distortion type.
- From Tables 4 and 5, it can be seen that the proposed SJSD metric has the lowest average standard deviation across different distortion types and levels. This

TABLE 4: Average standard deviation across models for all metrics for different levels of distortion.

	SAUC	SNSS	SEMD	SJSD	SSKLD
High	0.0510	0.33	0.0681	**0.0353**	0.2601
Medium	0.0257	0.1705	0.0190	**0.0144**	0.2259
Low	0.0281	0.2191	0.0115	**0.0069**	0.2588

TABLE 5: Average standard deviation across models for all metrics for different types of distortion.

	SAUC	SNSS	SEMD	SJSD	SSKLD
Blur	0.0768	0.4841	0.0810	**0.0409**	0.3945
JPEG	0.0253	0.1811	0.0323	**0.0197**	0.2291
Noise	0.0492	0.3128	0.0558	**0.0314**	0.2698

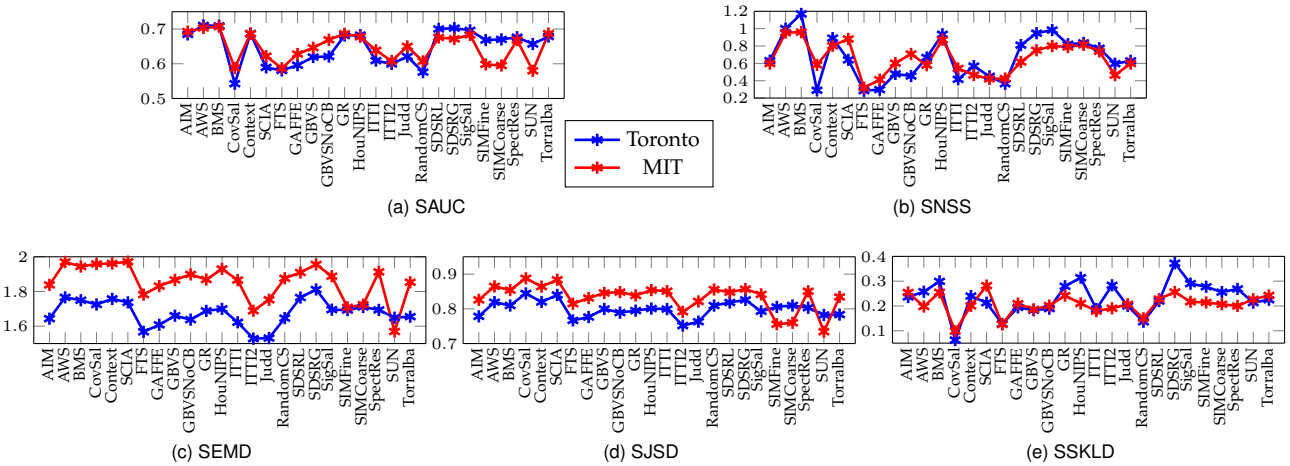

Fig. 12: Metric scores for the Toronto [28] and MIT [34] databases.

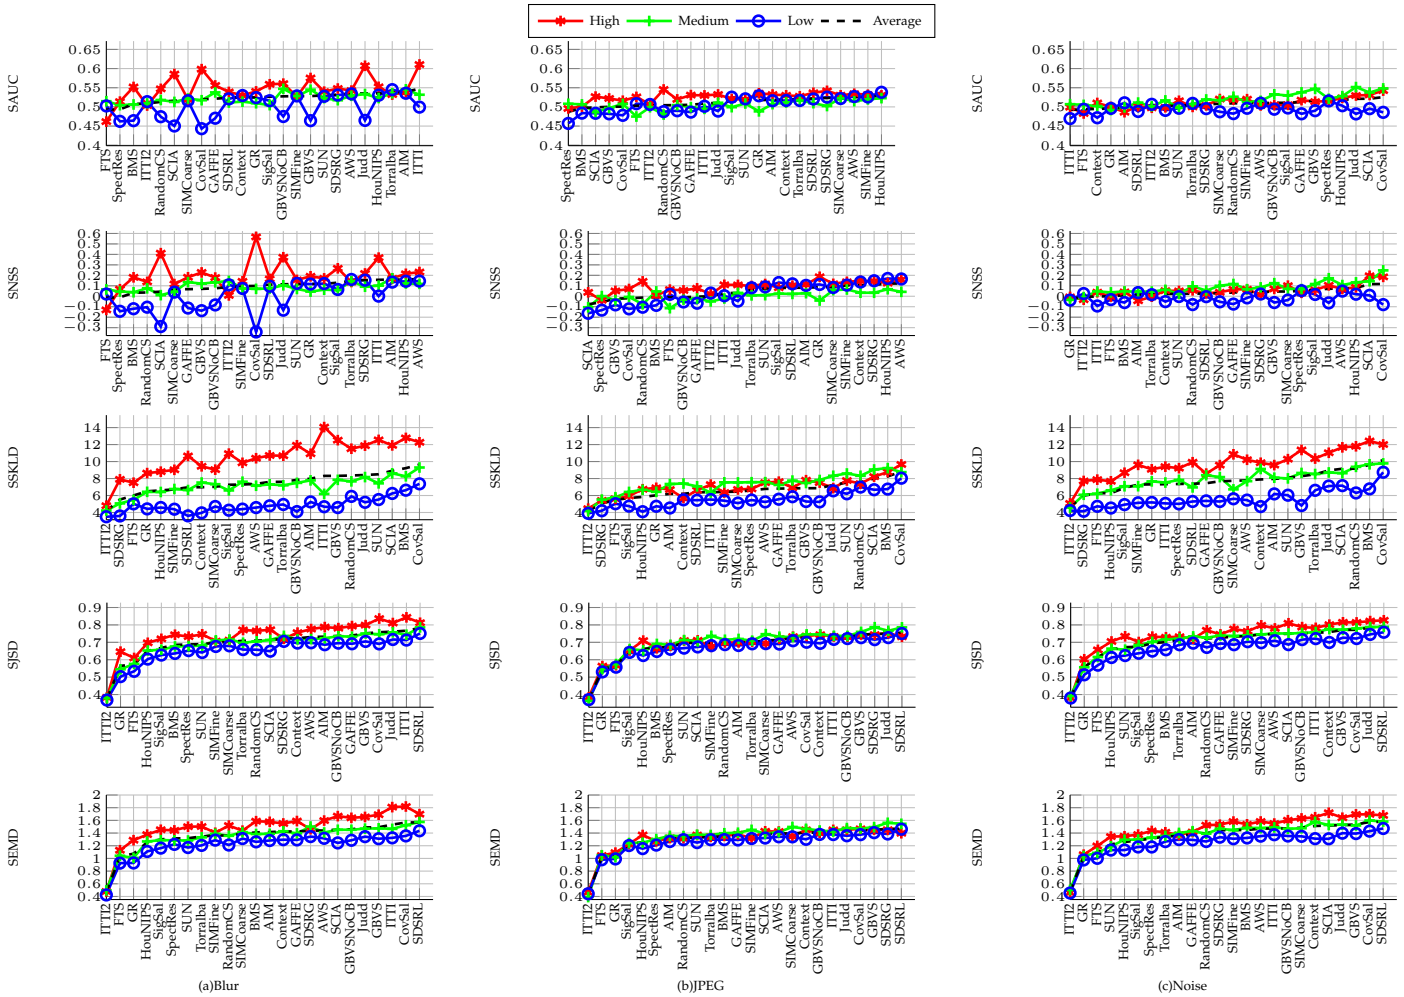

Fig. 13: SAUC, SNSS, SSKLD, SJSD, and SEMD scores for the TUD database for different distortion types:(a) Blur, (b) JPEG and (c) Noise. Blue, green and red lines correspond to low, medium and high levels of distortion, respectively, and the black dotted line represents the average over all distortion levels according to which the models have been sorted.

shows that the proposed SJSD metric is less sensitive (and thus more resilient) to the changes in the amount and type of distortion in the sense that it exhibits a more consistent and stable performance across different distortion types and levels as compared to the existing SAUC metric.

4.2.3 Variation in performance with varying distortion levels for the Judd Low Resolution Database

The Judd low resolution database has only one type of distortion, i.e., blur, but it has 8 different levels of distortions that correspond to the 8 different resolutions (512 to 4 in factors of 2). The images are also divided into three categories: easy, medium and hard, based on their complexity. We report the variation in performance for the three different image categories for only 3 out of the 8 resolutions for ease of presentation. The 3 resolutions chosen are 64, 128 and 256 which correspond to high, medium and low levels of blur, respectively. Figure 14 shows the variation in performance in terms of the SAUC, SNSS, SSKLD, SJSD and SEMD metrics, respectively, for the 3 different distortion levels. From these figures we observe the following:

- For the SSKLD and SJSD metrics, most of the models perform better for high and medium levels of distortions than that for the low or near-zero distortion case over all the image categories. For the SAUC metric, the performance at the high and medium levels of distortions is greater than that for the low level for most models over the hard and medium image categories. For the SEMD and SNSS metrics, no trend in scores is observed.
- For all the metrics, the average scores over all blur levels over all the models decrease with increase in image complexity. This is seen in the fall of the curve representing average scores (marked by a black dotted line in Figure 14) as the image complexity goes from easy to hard. This is to be expected as increase in image complexity increases the clutter in the scene thus making it harder for the VA models to distinguish the salient locations from the non-salient background.

5 Conclusion

In this paper, noticeable shortcomings of existing VA performance evaluation metrics are studied and four new performance metrics that remedy these shortcomings are proposed. Also, a comprehensive analysis of several state-of-the-art VA models is performed over two databases with different distortion types and several levels of distortions in addition to the distortion-free case. Along with the widely used SAUC metric, the performance is evaluated in terms of the four proposed shuffled metrics, SNSS, SSKLD, SJSD and SEMD. For the distortion-free case, it is shown that the proposed metrics achieve a consistent performance in terms of model rankings across different datasets. Furthermore, it is shown that the proposed SJSD metric is more robust than the other metrics, including the widely used SAUC metric, in the presence of various levels and types of distortions in the sense that it exhibits the most consistent performance across distortion levels and distortion types. It is found that the performance of the VA models varies with both the type and level of distortion. Also, it is found that, in general, high and medium levels of distortion improve the performance of the VA models relative to the low or near zero-distortion case. Finally, the VA models' performance is also found to decrease with an increase in image complexity.

References

[1] A. Borji and L. Itti, "State-of-the-art in visual attention modeling," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 35, no. 1, pp. 185–207, Jan 2013.

[2] A. Borji, D. Sihite, and L. Itti, "Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study," *IEEE Trans. Image Processing*, vol. 22, no. 1, pp. 55–69, Jan 2013.

[3] T. Judd, F. Durand, and A. Torralba, "A benchmark of computational models of saliency to predict human fixations," *MIT Tech Report*, 2012.

[4] D. Walther and C. Koch, "Modeling attention to salient proto-objects," *Neural networks*, vol. 19, no. 9, pp. 1395–1407, 2006.

[5] S. Frintrop, *VOCUS: A visual attention system for object detection and goal-directed search*. Springer, 2006.

[6] A. Mishra, Y. Aloimonos, and C. L. Fah, "Active segmentation with fixation," in *12th IEEE Int'l Conf. Computer Vision*, Sept 2009, pp. 468–475.

[7] A. Maki, P. Nordlund, and J.-O. Eklundh, "Attentional scene segmentation: integrating depth and motion," *Computer Vision and Image Understanding*, vol. 78, no. 3, pp. 351–373, 2000.

[8] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Processing*, vol. 13, no. 4, pp. 600–612, April 2004.

[9] H. Sheikh and A. Bovik, "Image information and visual quality," *IEEE Trans. Image Processing*, vol. 15, no. 2, pp. 430–444, Feb 2006.

[10] Z. Wang, E. Simoncelli, and A. Bovik, "Multiscale structural similarity for image quality assessment," in *the Thirty-Seventh Asilomar Conf. Signals, Systems and Computers*, vol. 2, Nov 2003, pp. 1398–1402.

[11] U. Engelke, H. Kaprykowsky, H.-J. Zepernick, and P. Ndjiki-Nya, "Visual attention in quality assessment," *IEEE Signal Processing Magazine*, vol. 28, no. 6, pp. 50–59, Nov. 2011.

[12] A. Ninassi, O. Le Meur, P. Le Callet, and D. Barbba, "Does where you gaze on an image affect your perception of quality? applying visual attention to image quality metric," in *IEEE Int'l Conf. Image Processing*, vol. 2, Sept 2007, pp. 169–172.

[13] E. Larson, C. Vu, and D. Chandler, "Can visual fixation patterns improve image fidelity assessment?" in *15th IEEE Int'l Conf. Image Processing*, Oct 2008, pp. 2572–2575.

[14] H. Liu and I. Heynderickx, "Visual attention in objective image quality assessment: Based on eye-tracking data," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 21, no. 7, pp. 971–982, Jul. 2011.

[15] R. Barland and A. Saadane, "Blind quality metric using a perceptual importance map for jpeg-20000 compressed images," *IEEE Int'l Conf. Image Processing*, pp. 2941–2944, 2006.

[16] A. K. Moorthy and A. C. Bovik, "Visual Importance Pooling for Image Quality Assessment," *IEEE J. Selected Topics in Signal Processing*, vol. 3, no. 2, pp. 193–201, Apr. 2009.

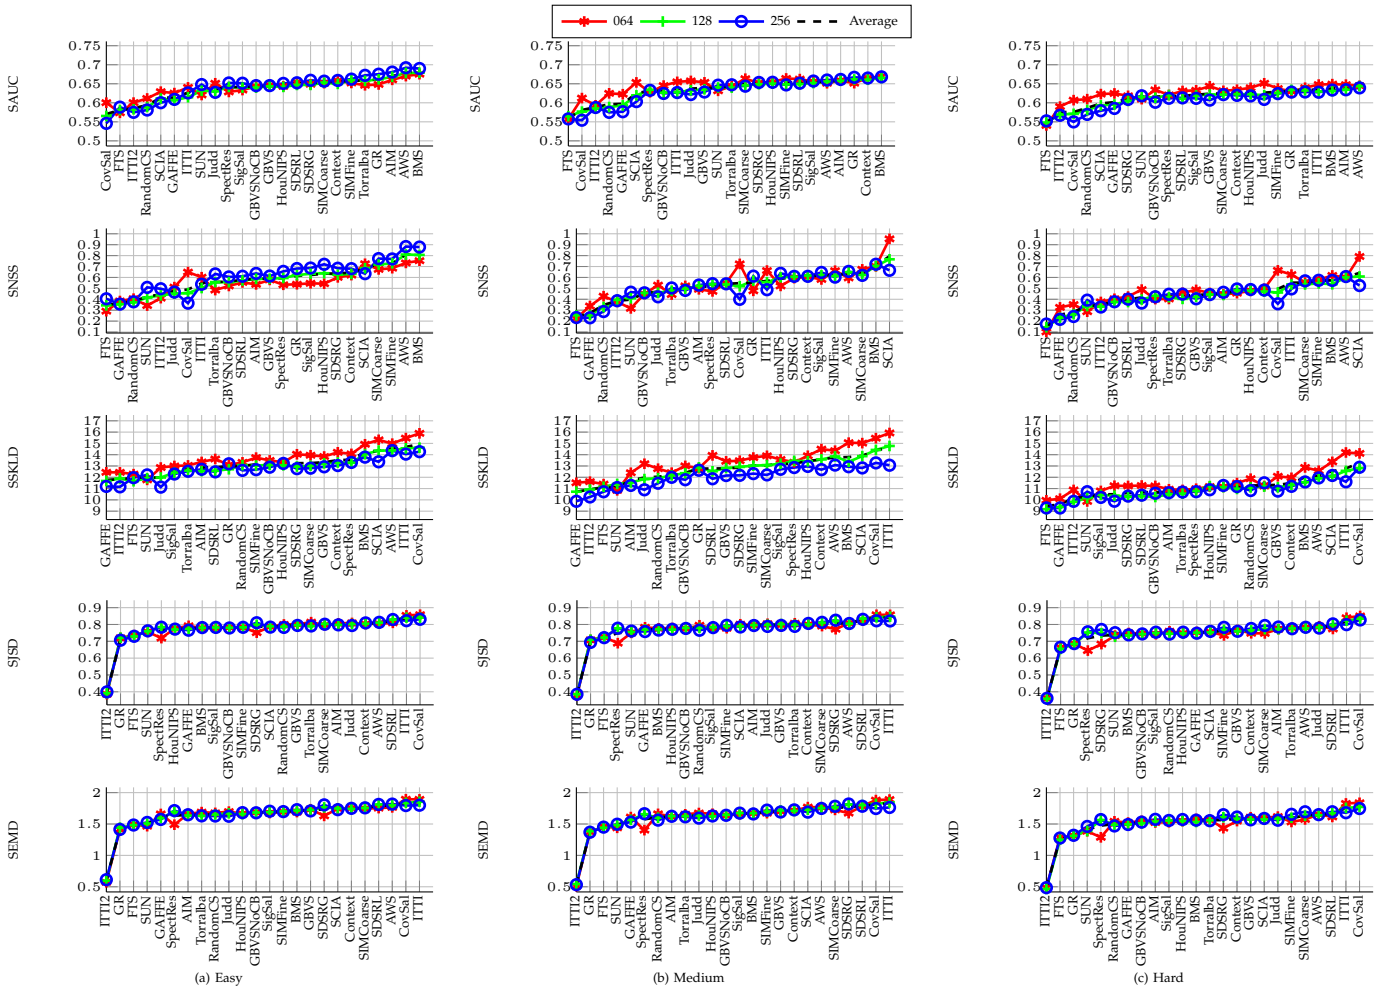

Fig. 14: SAUC, SNSS, SSKLD, SJSD, and SEMD scores for the Judd Low Resolution database for different image complexity categories:(a) Easy, (b) Medium and (c) Hard. Blue, green and red lines correspond to low, medium and high levels of blur, respectively, and the black dotted line represents the average over all distortion levels according to which the models have been sorted.

[17] N. G. Sadaka, L. J. Karam, R. Ferzli, and G. P. Abousleman, "A no-reference perceptual image sharpness metric based on saliency-weighted foveal pooling," in *IEEE Int'l Conf. Image Processing*, Oct. 2008, pp. 369–372.

[18] I. Gkioulekas, G. Evangelopoulos, and P. Maragos, "Spatial bayesian surprise for image saliency and quality assessment," in *IEEE Int'l Conf. Image Processing*, Sep. 2010, pp. 1081–1084.

[19] A. K. Moorthy and A. C. Bovik, "Visual importance pooling for image quality assessment," *IEEE J. Selected Topics in Signal Processing*, vol. 3, no. 2, pp. 193–201, Apr. 2009.

[20] H. Liu and I. Heynderickx, "Visual attention in objective image quality assessment: based on eye-tracking data," *IEEE Trans. Circuits and Systems for Video Technology*, vol. 21, no. 7, pp. 971–982, Jul. 2011.

[21] J. Redi, H. Liu, R. Zunino, and I. Heynderickx, "Interactions of visual attention and quality perception," *Proc. SPIE, Human Vision and Electronic Imaging XVI*, vol. 7865, no. 1, pp. 1–11, Feb. 2011.

[22] C. T. Vu, E. C. Larson, and D. M. Chandler, "Visual fixation patterns when judging image quality: Effects of distortion type, amount, and subject experience," *IEEE Southwest Symposium on Image Analysis and Interpretation, 2008*, pp. 73–76, Mar. 2008.

[23] A. L. Yarbus, "Eye movements during perception of complex objects," in *Eye Movements and Vision*. Springer, 1967, pp. 171–211.

[24] S. Mathe and C. Sminchisescu, "Action from still image dataset and inverse optimal control to learn task specific visual scanpaths," in *Advances in Neural Information Processing Systems*, C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, Eds., 2013, pp. 1923–1931.

[25] U. Rajashekar, I. van der Linde, A. C. Bovik, and L. K. Cormack, "GAFFE: a Gaze-Attentive fixation finding engine," *IEEE Trans. Image Processing*, vol. 17, no. 4, pp. 564–573, Apr. 2008.

[26] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 20, no. 11, pp. 1254–1259, Nov 1998.

[27] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," *Advances in Neural Information Processing Systems*, vol. 19, pp. 545–552, 2007.

[28] N. Bruce and J. Tsotsos, "Saliency, attention, and visual search: An information theoretic approach," *J. Vision*, vol. 9, no. 3, pp. 1–24, Mar. 2009.

[29] X. Hou and L. Zhang, "Dynamic visual attention: searching for coding length increments," in *Advances in Neural Information Processing Systems*, 2008, pp. 681–688.

[30] M. Mancas, C. Mancas-thillou, B. Gosselin, and B. Macq, "A rarity-based visual attention map-application to texture description," in *IEEE Int'l Conf. Image Processing*, Oct. 2006, pp. 445 –448.

[31] H. J. Seo and P. Milanfar, "Static and space-time visual saliency detection by self-resemblance," *J. Vision*, vol. 9, no. 12, p. 15, Nov. 2009.

[32] L. Zhang, M. H. Tong, T. K. Marks, H. Shan, and G. W. Cot-

trell, "SUN: a bayesian framework for saliency using natural statistics," *J. Vision*, vol. 8, no. 7, pp. 1 – 20, Dec. 2008.
[33] A. Oliva and A. Torralba, "Modeling the shape of the scene: A holistic representation of the spatial envelope," *Int'l J. Computer Vision*, vol. 42, no. 3, pp. 145–175, May 2001.
[34] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *12th IEEE Int'l Conf. Computer Vision*, Oct. 2009, pp. 2106–2113.
[35] H. Tavakoli, E. Rahtu, and J. Heikkila, "Fast and efficient saliency detection using sparse sampling and kernel density estimation," in *Scandinavian Conf. Image Analysis*. Springer Berlin Heidelberg, 2011, pp. 666–675.
[36] R. Achanta, S. Hemami, F. Estrada, and S. Susstrunk, "Frequency-tuned salient region detection," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Jun. 2009, pp. 1597–1604.
[37] X. Hou, J. Harel, and C. Koch, "Image signature: Highlighting sparse salient regions," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 34, pp. 194–201, Jul. 2011.
[38] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2007, pp. 1–8.
[39] A. Garcia-Diaz, V. Leborán, X. R. Fdez-Vidal, and X. M. Pardo, "On the relationship between optical variability, visual saliency, and eye fixations: a computational approach," *J. Vision*, vol. 12, no. 6, pp. 1–22, Jan. 2012.
[40] J. Zhang and S. Sclaroff, "Saliency detection: A boolean map approach," in *Proc. of the IEEE Int'l Conf. Computer Vision (ICCV)*, Dec. 2013, pp. 153–160.
[41] S. Goferman, L. Zelnik-Manor, and A. Tal, "Context-aware saliency detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 34, no. 10, pp. 1915–1926, 2012.
[42] E. Erdem and A. Erdem, "Visual saliency estimation by nonlinearly integrating features using region covariances," *J. Vision*, vol. 13, no. 4, p. 11, Mar. 2013, PMID: 23509407.
[43] T. N. Vikram, M. Tscherepanow, and B. Wrede, "A saliency map based on sampling an image into random rectangular regions of interest," *Pattern Recognition*, vol. 45, no. 9, pp. 3114 – 3124, 2012.
[44] R. Margolin, L. Zelnik-Manor, and A. Tal, "Saliency for image manipulation," *The Visual Computer*, vol. 29, no. 5, pp. 381–392, May 2013.
[45] C. Koch and S. Ullman, "Shifts in selective visual attention: Towards the underlying neural circuitry," in *Matters of Intelligence*. Springer, 1987, pp. 115–141.
[46] A. M. Treisman and G. Gelade, "A feature-integration theory of attention," *Cognitive Psychology*, vol. 12, no. 1, pp. 97–136, 1980.
[47] C. E. Shannon, "A mathematical theory of communication," *ACM SIGMOBILE Mobile Computing and Communications Rev.*, vol. 5, no. 1, pp. 3–55, 2001.
[48] A. Torralba, "Modeling global scene factors in attention," *J. the Optical Society of America. A, Optics, Image Science, and Vision*, vol. 20, no. 7, pp. 1407–1418, 2003.
[49] T. Judd, F. Durand, and A. Torralba, "A benchmark of computational models of saliency to predict human fixations," in *MIT Technical Report*, 2012.
[50] Y. Rubner, C. Tomasi, and L. J. Guibas, "The earth mover's distance as a metric for image retrieval," *Int'l J. Computer Vision*, vol. 40, no. 2, pp. 99–121, Nov. 2000.
[51] O. Pele and M. Werman, "A linear time histogram metric for improved sift matching," in *10th European Conf. Computer Vision (ECCV)*. Springer, Oct 2008, pp. 495–508.
[52] O. Pele and M. Werman, "Fast and robust earth mover's distances," in *12th Int'l Conf. Computer Vision (ICCV)*, Sept 2009, pp. 460–467.
[53] B. W. Tatler, R. J. Baddeley, and I. D. Gilchrist, "Visual correlates of fixation selection: Effects of scale and time," *Vision Research*, vol. 45, pp. 643–659, 2005.
[54] A. Mittal, A. Moorthy, and A. Bovik, "Automatic prediction of saliency on JPEG distorted images," in *2011 Third Int'l Workshop on Quality of Multimedia Experience (QoMEX)*, September 2011, pp. 195 –200.
[55] R. J. Peters, A. Iyer, L. Itti, and C. Koch, "Components of bottom-up gaze allocation in natural images." *Vision research*, vol. 45, no. 18, pp. 2397–416, Aug. 2005.
[56] R. Margolin, L. Zelnik-Manor, and A. Tal, "How to evaluate foreground maps," in *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2014, pp. 248–255.
[57] S. Kullback and R. A. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, pp. 79–86, 1951.
[58] N. Riche, M. Duvinage, M. Mancas, B. Gosselin, and T. Dutoit, "Saliency and human fixations: State-of-the-art and study of comparison metrics," in *IEEE Int'l Conf. Computer Vision (ICCV)*, Dec 2013, pp. 1153–1160.
[59] J. Lin, "Divergence measures based on the Shannon entropy," *IEEE Trans. Information Theory*, vol. 37, no. 1, pp. 145–151, 1991.
[60] D. M. Green, J. A. Swets *et al.*, *Signal Detection Theory and Psychophysics*. Wiley New York, 1966, vol. 1.
[61] D. Endres and J. Schindelin, "A new metric for probability distributions," *IEEE Trans. Information Theory*, vol. 49, no. 7, pp. 1858–1860, July 2003.
[62] H. Liu and I. Heynderickx, "Studying the added value of visual attention in objective image quality metrics based on eye movement data," in *16th IEEE Int'l Conf. Image Processing (ICIP)*, Nov. 2009, pp. 3097 –3100.
[63] T. Judd, F. Durand, and A. Torralba, "Fixations on low-resolution images," *J. Vision*, vol. 11, pp. 1–20, 2011.

Milind S. Gide received the M.S. degree in Electrical Engineering from the University of Houston, Houston, TX, USA in 2009 . He is currently a Ph.D. student in the Image, Video, & Usability lab and a teaching assistant in the School of Electrical, Computer & Energy Engineering at Arizona State University. His research interests include visual attention, image quality and texture-based video coding. During Summer 2011, he worked as an intern at Sharp Labs of America, Camas, Washington in the area of texture-based video coding.

Samuel F. Dodge received the B.S.E and M.S. both in Electrical Engineering from Arizona State University, Tempe AZ USA in 2011 and 2014, respectively. In 2013 he was a visiting researcher at Peking University in Beijing, China as part of the NSF EAPSI program. He has also held internships at Intel in Chandler, Arizona and at Sony in Kanagawa, Japan working on content-based image retrieval and biomedical image processing. He is currently a Ph.D. student in the Image, Video, & Usability Lab at Arizona State University, and a teaching assistant in the School of Electrical, Computer and Energy Engineering. He is a recipient of the ARCS fellowship in 2013 and 2014. His research interests include computational visual saliency and applications in computer vision and image processing.

Lina J. Karam (Fellow, IEEE) received the B.E. degree in computer and communications engineering from the American University of Beirut in 1989, and the M.S. and Ph.D. degrees in electrical engineering from the Georgia Institute of Technology in 1992 and 1995, respectively. She is a Full Professor in the School of Electrical, Computer & Energy Engineering at Arizona State University, where she directs the Image, Video, & Usabilty (IVU) Research Laboratory. Dr. Karam was awarded a U.S. National Science Foundation CAREER Award, a NASA Technical Innovation Award, the Intel Outstanding Researcher Award, and the IEEE Phoenix Section Outstanding Faculty Award. She served on several journal editorial boards, several conference organization committees, and several IEEE technical committees. She is currently serving as the General Chair of the 2016 IEEE International Conference on Image Processing. Professor Karam is a member of the IEEE Circuits and Systems Society's DSP Technical Committee, the IEEE Signal Processing Society's (SPS) IVMSP Technical Committee, and the IEEE SPS Board of Governors. She has over 200 technical publications and she is a co-inventor on a number of patents.